

Edge distribution in generalized graph products

Michael Langberg ^{*} Dan Vilenchik [†]

Abstract

Given a graph $G = (V, E)$, an integer k , and a function $f_G : V^k \times V^k \rightarrow \{0, 1\}$, the k^{th} graph product of G w.r.t f_G is the graph with vertex set V^k , and an edge between two vertices $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{y} = (y_1, \dots, y_k)$ iff $f_G(\mathbf{x}, \mathbf{y}) = 1$. Graph products are a basic combinatorial object, widely studied and used in different areas such as hardness of approximation, information theory, etc. We study graph products for functions f_G of the form $f_G(\mathbf{x}, \mathbf{y}) = 1$ iff there are at least t indices $i \in [k]$ s.t. $(x_i, y_i) \in E$, where $t \in [k]$ is a fixed parameter in f_G . This framework generalizes the well-known graph *tensor-product* (obtained for $t = k$) and the graph *or-product* (obtained for $t = 1$). The property that interests us is the edge distribution in such graphs. We show that if G has a spectral gap, then the number of edges connecting “large-enough” sets in G^k is “well-behaved”, namely, it is close to the expected value, had the sets been random. We extend our results to *bi-partite graph* products as well. For a bi-partite graph $G = (X, Y, E)$, the k^{th} bi-partite graph product of G w.r.t f_G is the bi-partite graph with vertex sets X^k and Y^k and edges between $\mathbf{x} \in X^k$ and $\mathbf{y} \in Y^k$ iff $f_G(\mathbf{x}, \mathbf{y}) = 1$. Finally, for both types of graph products, optimality is asserted using the “Converse to the Expander Mixing Lemma” obtained by Bilu and Linial in 2006. A byproduct of our proof technique is a new explicit construction of a family of co-spectral graphs.

^{*}Department of Mathematics and Computer Science, The Open University of Israel, 108 Ravutski St., Raanana 43107, Israel, email: mikel@openu.ac.il. Work supported in part by ISF grant 480/08.

[†]Dept. of Computer Science and Applied Math., The Weizmann Institute of Science, Rehovot 76100, Israel, email dan.vilenchik@weizmann.ac.il. Work done in part while at The Open University of Israel.

1 Introduction

Let $G = (V, E)$ be a given undirected graph and let k be an integer. For a function $f_G : V^k \times V^k \rightarrow \{0, 1\}$, the well-known graph product operation G^k is the graph with vertex set V^k , and an edge between $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{y} = (y_1, \dots, y_k)$ iff $f_G(\mathbf{x}, \mathbf{y}) = 1$. The graph product operation with parameter k is sometimes referred to as the k^{th} power of G .

There are several well-known graph products that have been studied in the literature. The tensor product, introduced by Alfred North Whitehead and Bertrand Russell in their 1912 *Principia Mathematica*, is defined using the function $f_G(\mathbf{x}, \mathbf{y}) = 1$ iff all indices i satisfy $(x_i, y_i) \in E$. This graph product appears in the literature also as the weak product, or the conjunction product. It is equivalent to the Kronecker product of graph adjacency matrices [16]. A famous standing conjecture related to tensor products is Hedetniemi's conjecture from 1966 [7], which states that the chromatic number of the tensor product of two graphs G_1, G_2 is at most the minimum between the chromatic numbers of G_1 and G_2 .

The graph or-product (also called the co-normal/disjunctive graph product) is defined using the function $f_G(\mathbf{x}, \mathbf{y}) = 1$ iff there exists at least one index i s.t. $(x_i, y_i) \in E$. A slightly different definition of the or-product was used by Garey and Johnson in their famous book [6] to prove the first hardness of approximation results for the graph coloring problem and the largest independent set problem (actually the book cover itself features a drawing of a graph product). In general, graph products are used as a tool in establishing hardness of approximation results, and specifically in amplifying existing hardness results.

Other types of graph products were studied as well, for example the Xor graph product in which $f_G(\mathbf{x}, \mathbf{y}) = 1$ iff an odd number of coordinates i satisfy $(x_i, y_i) \in E$. This product was studied in [13, 2]. One can generalize the definitions above by taking the product of k different graphs G_1, G_2, \dots, G_k , and $f : (V_1 \times \dots \times V_k)^2 \rightarrow \{0, 1\}$. In this paper we study the case where for every i , $G_i = G$, and we note that under suitable conditions on the G_i 's, our results can be stated in the general form.

1.1 Our Contribution and Techniques

In this paper we study two natural generalizations of the “or” and “tensor” products. Observe that the only difference in the definition of the two products is in the edge-threshold parameter t in f_G : at least one index (x_i, y_i) satisfying $(x_i, y_i) \in E$ for the or-product, and all k indices satisfying $(x_i, y_i) \in E$ for the tensor product. In this work we study graph products with an arbitrary threshold t . Formally, for an integer $t \in [k]$, let $f_G^{(t)}$ be the function $f_G^{(t)}(\mathbf{x}, \mathbf{y}) = 1$ iff there exist at least t indices i s.t. $(x_i, y_i) \in E$. Using this notation, the tensor product is obtained by taking $t = k$ and the or-product by $t = 1$. Our notion of graph product is similar in spirit to the special graph products known as the non-complete extended P -sum (NEPS) graph products, introduced for the first time by Cvetković in the early 1970's [4], and elaborated in [5]. One generalization of NEPS graph products given in [12] actually captures our notion of $f_G^{(t)}$ -graph product as a special case. More details are given in the technical sections.

The second generalization that we suggest is a *bi-partite* graph product with threshold t . Specifically, let $G = (X, Y, E)$ be a bi-partite graph, and define the k^{th} bi-partite power of G to be the graph with vertex sets $X^k \cup Y^k$ and edges (\mathbf{x}, \mathbf{y}) iff $f_G^{(t)}(\mathbf{x}, \mathbf{y}) = 1$.

In this work we consider only regular graphs G , and study the edge distribution in the k^{th} power of

G . A well-known edge discrepancy result for regular graphs is the *Expander Mixing Lemma* [1] (EML for short). The EML states that for a d -regular n -vertex graph G , with second largest eigenvalue (in absolute value) λ , and for any two sets of vertices S, T , the number of edges they span in G , $e(S, T) = \{(u, v) \in E \mid u \in S, v \in T\}$, satisfies

$$\left| e(S, T) - \frac{d|S||T|}{n} \right| \leq \lambda \sqrt{|S||T|}. \quad (1.1)$$

The quantity $d|S||T|/n$ can be thought of as the expected number of edges spanned by random sets S and T . The intuitive way to read the EML is that if $\lambda < d$, then the number of edges between “large enough” sets is well-behaved, namely it is close to the expected value. The EML as stated gives a trivial bound for d -regular bi-partite graphs G (since $\lambda(G) = |-d| = d$ as well). One can rather easily extend and restate the EML for d -regular bi-partite graphs (in which case λ in Eq. (1.1) refers to the third largest eigenvalue in absolute value, and a factor 2 multiplies $d|S||T|/n$):

$$\left| e(S, T) - \frac{d|S||T|}{n/2} \right| \leq \lambda \sqrt{|S||T|}. \quad (1.2)$$

Our study focuses on applying the Expander Mixing Lemma to our generalized graph products. The main challenge, after showing that G^k is regular and computing its degree, is to obtain an upper bound on λ . The spectrum of the tensor product, i.e. the case $t = k$, is well-understood, and has a simple clean characterization [11]. For $t < k$ one has to work harder to compute λ . The more challenging case is the bi-partite graph product. While previous results provide access to the spectrum of the non-bipartite graph product, no such result exists for the bi-partite product. The main technical achievement of this paper is giving an exact characterization of the spectrum of G^k in the bi-partite setting. This is obtained by embedding G^k in a certain larger graph, for which the spectrum can be analyzed, and then infer the spectrum of G^k from that graph. Let us remark that for both types of graph products, optimality results can be obtained from the “converse to the EML” result obtained in [3].

Another way of viewing our results is using the terminology of **(d, α) -jumbled graphs**, coined in [14], and used widely in the study of pseudo-random graphs (see the excellent survey in [10] for more details). A d -regular graph G is called (d, α) -jumbled if any two sets of vertices S, T satisfy Eq. (1.1), with α replacing λ . We compute the value α , s.t. G^k is (d', α) -jumbled (d' is the degree of G^k). We do that for both types of graph products, and for every parameter t .

Finally let us mention a byproduct of our proof technique, which may be of self interest. Recall that the spectrum of a graph G is the multiset of eigenvalues of its adjacency matrix. Two graphs are called **cospectral** if they have the same spectrum. Clearly isomoprhic graphs are cospectral, but the interesting question is to identify pairs of graphs which are not isomorphic, yet cospectral. There are numerous constructions of such pairs. Seidel switching, introduced in [15], is a quite general method to construct such pairs, and was used in [9] for example. In this paper we give a new construction of co-spectral graphs. This construction is natural and simple to describe, and uses a different approach, which is closer in flavor to the construction in [8]. Instead of local switching, our technique can be viewed as a family of blow-ups of a d -regular bi-partite graph G that produce a family of co-spectral graphs.

The remainder of the paper is organized as follows. In Section 2 we formally state our results. All the technical details and proofs follow in the subsequent sections.

2 Main Results

We start with the formal definition of the graph product. For sake of precision, we replace the notation “ G^k ”, which we used in the Introduction, with a more detailed and explicit notation.

Definition 2.1. For a simple graph $G = (V, E)$, define the **graph product** $GP_{k,t}(G)$ to be the graph with vertex set V^k , and an edge $(x_1, \dots, x_k) \sim (y_1, \dots, y_k)$ if there are at least t pairs (x_i, y_i) that share an edge in G .

Definition 2.2. If $G = (X, Y, E)$ is a simple bipartite graph, define the **bi-partite graph product** $BGP_{k,t}(G)$ to be the graph with vertex set (X^k, Y^k) , and an edge between $(x_1, \dots, x_k) \in X^k$ and $(y_1, \dots, y_k) \in Y^k$ if there are at least t pairs (x_i, y_i) that share an edge in G .

Let us first establish the easy fact that both graph products are regular graphs, and state their degrees.

Proposition 2.3. Let G be an n -vertex d -regular graph, and G' be an n -vertex bi-partite d -regular graph. Then $GP_{k,t}(G)$ is d_1 -regular and $BGP_{k,t}(G')$ is d_2 -regular with d_1 and d_2 given by:

$$d_1 = \sum_{t'=t}^k \binom{k}{t'} d^{t'} (n-d)^{k-t'}, \quad d_2 = \sum_{t'=t}^k \binom{k}{t'} d^{t'} \left(\frac{n}{2} - d\right)^{k-t'}. \quad (2.1)$$

The proof of Proposition 2.3 is rather straightforward and is given in Section 6.

Before we state our main results, we introduce some notations: We let $e(S, T)$ stand for the number of edges connecting two sets of vertices S and T (when the underlying graph is clear from context). For a graph G , we let $A(G)$ be its adjacency matrix. Somewhat abusing notations, when we refer to graph eigenvalues, we mean the eigenvalues of its adjacency matrix. For a graph G with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$, let $\lambda(G) = \max\{\lambda_2, |\lambda_n|\}$. If G is bi-partite then $\lambda(G) = \max\{\lambda_2, |\lambda_{n-1}|\}$ (the reason is that $\lambda_1 = |\lambda_n| = d$ if G is bi-partite). For a D -regular N -vertex graph, and for any two sets of vertices S, T , we denote by $\mu_{S,T}$ the expected number of edges that connect S and T , had the sets been chosen uniformly at random. The following equations give this quantity for the regular and bi-partite case:

$$\mu_{S,T} = \frac{D|S||T|}{N}, \quad \mu_{S,T} = \frac{D|S||T|}{N/2} = \frac{2D|S||T|}{N}. \quad (2.2)$$

Our first theorem establishes a tight upper bound on the second largest eigenvalue of $GP_{k,t}(G)$.

Theorem 2.4. Let $GP_{k,t}(G)$ be the graph product of an n -vertex d -regular graph $G = (V, E)$ with $\lambda = \lambda(G)$. Let $\Lambda = \lambda(GP_{k,t}(G))$. If $d \leq (n-1)/2$ then

$$\frac{1}{3^k} \cdot \frac{\lambda}{d} \cdot d_1 \leq \Lambda \leq \frac{\lambda}{d} \cdot d_1$$

The following result about the edge-distribution in the non-bipartite graph product is an immediate corollary of Theorem 2.4 and the EML (Eq. (1.1)):

Corollary 2.5. Under the conditions of Theorem 2.4, for any two sets of vertices S, T in $GP_{k,t}(G)$

$$|e(S, T) - \mu_{S,T}| \leq \frac{\lambda}{d} \cdot d_1 \sqrt{|S||T|}. \quad (2.3)$$

The next theorem is the analogue of Theorem 2.4 for the bi-partite setting. While the adjacency matrix of $GP_{k,t}(G)$ has an explicit closed expression (based on tensor products), this is not the case for $BGP_{k,t}(G)$. Therefore the task of estimating Λ is more complicated.

Recall that for a bi-partite graph G , $\lambda(G)$ is the third largest eigenvalue in absolute value.

Theorem 2.6. *Let $BGP_{k,t}(G)$ be the bi-partite graph product of an n -vertex bi-partite d -regular graph $G = (X, Y, E)$ with $\lambda = \lambda(G)$. Let Λ be the third largest eigenvalue in absolute value of $BGP_{k,t}(G)$. If $d \leq n/4$ then*

$$\frac{1}{2^k} \cdot \frac{\lambda}{d} \cdot d_2 \leq \Lambda \leq \frac{\lambda}{d} \cdot d_2.$$

Consequently, together with the EML for bi-partite graphs given in Eq. (1.2), we get the following edge-distribution result for $BGP_{k,t}(G)$:

Corollary 2.7. *Under the conditions of Theorem 2.6, for any two sets of vertices $S \subseteq X^k$ and $T \subseteq Y^k$,*

$$|e(S, T) - \mu_{S,T}| \leq \frac{\lambda}{d} \cdot d_2 \sqrt{|S||T|}. \quad (2.4)$$

A few remarks concerning Theorems 2.4 and 2.6 are in place.

1. The upper bound on the discrepancy in Corollaries 2.5 and 2.7 is tight (up to $\log d$ -factors, and factors that depend on k). This follows from the converse to the EML result proven in [3], together with our lower bound on Λ in Theorems 2.4 and 2.6. The converse to the EML reads as follow: Let H be an N -vertex D -regular graph. If for any S, T with $S \cap T = \emptyset$ $|e(S, T) - \mu_{S,T}| \leq \alpha \sqrt{|S||T|}$, then all but the largest eigenvalue of H are bounded in absolute value by $O(\alpha(1 + \log(D/\alpha)))$. A similar result is derived in [3] for a bi-partite regular graph, in which case the same upper bound is obtained on all but the two largest eigenvalues.
2. We can restate Theorems 2.4 and 2.6 in the jumbled-graph terminology [14], namely $GP_{k,t}(G)$ is $(d_1, \frac{\lambda}{d} \cdot d_1)$ -jumbled, and $BGP_{k,t}(G)$ is $(d_2, \frac{\lambda}{d} \cdot d_2)$ -jumbled.
3. If $t = k$, namely $GP_{k,t}(G)$ is the tensor product, then it is known that $\lambda(GP_{k,t}(G)) = \lambda d^{k-1}$, $d_1 = d^k$, or in other words $\lambda(GP_{k,t}(G)) = (\frac{\lambda}{d}) d_1$. This exactly coincides with our upper bound on Λ given in Theorem 2.4. In fact, we obtain an exact expression for Λ , given in Eq. (3.3), and for $t = k$, it reads exactly $\Lambda = \lambda d^{k-1}$.
4. The assumptions on d in the statement of both theorems is made to enable a succinct description for the estimates on Λ , and as we mentioned previously, Eq. (3.3) and (4.3) spell an exact expression for Λ .

Finally, let us discuss an **application** of Theorem 2.4. Let G be a d -regular expander graph, say $d \leq n/2$ and $\lambda \leq 2\sqrt{d}$. In this case $\lambda < n - d - 1$ (a requirement in Theorem 2.4). Consider two sets of vertices S, T in $GP_{k,t}(G)$, and for simplicity assume they have the same size $|S| = |T| = \xi n^k$ for some $\xi \in [0, 1]$. For our choice of parameters, the right hand side of Eq. (2.3) is at most

$$\frac{\lambda}{d} \cdot d_1 \sqrt{|S||T|} \leq \frac{2}{\sqrt{d}} \cdot d_1 |S| = \frac{2}{\sqrt{d}} \cdot \frac{d_1 |S||T|}{n^k} \cdot \frac{n^k}{|T|} = \frac{2}{\xi \sqrt{d}} \cdot \mu_{S,T}.$$

If, say, $\xi \geq \log d / \sqrt{d}$, then the latter together with Eq. (2.3) imply that

$$e(S, T) = (1 \pm \varepsilon) \frac{d_1 |S||T|}{n^k},$$

where $\varepsilon \rightarrow 0$ as $d \rightarrow \infty$. In words, all “sufficiently large” sets of vertices have the “right” number of edges between them.

2.1 A family of co-spectral graphs

In Definition 2.2 we specified the bi-partite product $BGP_{k,t}(G)$ of a graph $G = (X, Y, E)$ as a graph over the vertex set $X^k \cup Y^k$. Namely, each vertex in $BGP_{k,t}(G)$ is either a k -tuple consisting of elements from X , or a k -tuple consisting of elements from Y . In our proofs to follow, we study additional forms of bipartite graph products for G in which vertices in the product graph are k -tuples which may consist of a mixture of elements from both X and Y .

Formally, let \mathcal{X} and \mathcal{Y} be symbols representing the sets X and Y respectively. We define the concept of a **template** $\tau \in \{\mathcal{X}, \mathcal{Y}\}^k$, which is a string of length k consisting of \mathcal{X} and \mathcal{Y} symbols.

A k -tuple of vertices $\mathbf{a} = (a_1, \dots, a_k) \in (X \cup Y)^k$ is said to have template $\tau = \tau_1, \dots, \tau_k$ iff $a_i \in X$ whenever $\tau_i = \mathcal{X}$ (and equivalently, $a_i \in Y$ whenever $\tau_i = \mathcal{Y}$). For a template τ , we define its complement template τ^c by switching \mathcal{X} to \mathcal{Y} and vice versa. For example, all vertices in $BGP_{k,t}(G)$ have either the template $\tau = \mathcal{X}^k$ or its complement $\tau^c = \mathcal{Y}^k$.

For a given template τ , denote the k -tuples in $(X \cup Y)^k$ with template τ by V_τ . Now, given a bi-partite graph $G = (X \cup Y, E)$, and a template τ , we denote by $BGP_{k,t,\tau}$ the bi-partite graph product with vertex set $V_\tau \cup V_{\tau^c}$, in which two vertices $\mathbf{a} = (a_1, \dots, a_k)$ and $\mathbf{b} = (b_1, \dots, b_k)$ are connected iff at least t pairs (a_i, b_i) are edges in G . In this work, we study the graphs $BGP_{k,t,\tau}$ and use them to prove our main results. On the way, we prove that

Theorem 2.8. *Let $G = (X \cup Y, E)$ be a d -regular graph. Then for every $k \geq 1$ and any two templates τ, τ' , the graphs $BGP_{k,t,\tau}$ and $BGP_{k,t,\tau'}$ are co-spectral.*

The proof is self contained and is given in full in Section 5.

3 Proof of Theorem 2.4

Let $G = (V, E)$ be a d -regular n -vertex graph with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, and corresponding eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n$. Similarly let $\Lambda_1, \Lambda_2, \dots, \Lambda_{n^k}$ be the eigenvalues of $GP_{k,t}(G)$, with corresponding eigenvectors $\mathbf{w}_1, \dots, \mathbf{w}_{n^k}$. In a slightly different setting and using different terminology, $GP_{k,t}(G)$ was studied in [12] and the following result regarding its spectrum was obtained (we rephrase that result to match our terminology).

Theorem 3.1. [12, Thm 3] *Let G be an n -vertex d -regular graph with eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, and corresponding eigenvectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$. Define $\lambda_1^* = n - 1 - d$, and $\lambda_i^* = -1 - \lambda_i$ for all $i \geq 2$. Let \mathcal{B} be the set of all vectors in $\{-1, 0, 1\}^k$ having at least t 1-entries. Then $GP_{k,t}(G)$ has a spectrum consisting of eigenvalues*

$$\Lambda_{i_1 i_2 \dots i_k} = \sum_{\mathbf{b} \in \mathcal{B}} \prod_{j=1}^k \left(\frac{1 + \mathbf{b}_j}{2} (\lambda_{i_j})^{|\mathbf{b}_j|} + \frac{1 - \mathbf{b}_j}{2} (\lambda_{i_j}^*)^{|\mathbf{b}_j|} \right), \quad i_1, \dots, i_k \in \{1, \dots, n\}. \quad (3.1)$$

The corresponding eigenvector is $\mathbf{w}_{i_1 i_2 \dots i_k} = \mathbf{u}_{i_1} \otimes \mathbf{u}_{i_2} \dots \otimes \mathbf{u}_{i_k}$.

Our proof outline is as follows: we first show that $\Lambda_{11\dots 1}$ is the largest eigenvalue of $GP_{k,t}(G)$, we then upper bound $\Lambda = \lambda(GP_{k,t}(G))$, which is the maximum over all $\Lambda_{i_1 i_2 \dots i_k}$'s in which not all i_j are equal to 1.

To prove that $\Lambda_{11\dots 1}$ is the largest eigenvalue of $GP_{k,t}(G)$ it suffices to show that it equals d_1 , as it is well-known that the largest eigenvalue of a D -regular graph is simply D . Using Eq. (3.1) and the definition of $\lambda_1 = d$ and $\lambda_1^* = n - d - 1$, we get the following expression for $\Lambda_{11\dots 1}$:

$$\Lambda_{11\dots 1} = \sum_{\mathbf{b} \in \mathcal{B}} \prod_{j=1}^k \left(\frac{1+\mathbf{b}_j}{2} d^{|\mathbf{b}_j|} + \frac{1-\mathbf{b}_j}{2} (n-d-1)^{|\mathbf{b}_j|} \right). \quad (3.2)$$

Next we show that Eq. (3.2) equals d_1 . This calculation was carried in [12], and we present it here for the sake of completeness. Fix a vertex $\mathbf{x} = (x_1, \dots, x_k)$, and let us compute the number of vertices (y_1, \dots, y_k) adjacent to \mathbf{x} with respect to a fixed $\mathbf{b} \in \mathcal{B}$ according to the following rules: if $\mathbf{b}_i = 1$ then (x_i, y_i) is an edge in E (in this case there are d choices for y_i); if $\mathbf{b}_i = 0$ then $x_i = y_i$ (one choice for y_i); if $\mathbf{b}_i = -1$ then $x_i \neq y_i$ and $(x_i, y_i) \notin E$ ($n - d - 1$ choices for y_i). This is reflected exactly in the product in Eq. (3.2). Since every vector \mathbf{b} specifies a different configuration of neighbors/non-neighbors of \mathbf{x} , and \mathcal{B} contains all possible configurations, we obtain $\Lambda_{11\dots 1} = d_1$.

Finally we compute Λ . Using our last insight about $\Lambda_{11\dots 1}$, the largest eigenvalue of $GP_{k,t}(G)$, we may conclude that Λ is the maximum over all $\Lambda_{i_1 i_2 \dots i_k}$ in which at least one i_j satisfies $i_j > 1$. Intuitively, for every $i_j > 1$, we loose a factor proportional to λ_{i_j}/d in the product of Eq. (3.1). Therefore, we expect Λ to correspond to $\Lambda_{i_1 i_2 \dots i_k}$ where exactly one $i_j > 1$. Let us start by computing the value of such eigenvalues, and then explain why indeed they give Λ . By symmetry it doesn't matter which $i_j > 1$, so let's assume $i_1 > 1$.

Denote $\Psi_{s,\mathbf{b}} = \prod_{j=s}^{\#\mathbf{b}} \left(\frac{1+\mathbf{b}_j}{2} d^{|\mathbf{b}_j|} + \frac{1-\mathbf{b}_j}{2} (n-d-1)^{|\mathbf{b}_j|} \right)$ (here $\#\mathbf{b}$ is the length of the vector \mathbf{b} which, as seen below, will change throughout the proof), and let us compute $\Lambda_{i11\dots 1}$ for some $i > 1$. Using these notations and Eq. (3.1) we obtain

$$\Lambda_{i11\dots 1} = \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} \lambda_i \Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=0} \Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=-1} (-\lambda_i - 1) \Psi_{2,\mathbf{b}}.$$

Define $\mathcal{C}_r = \{\mathbf{c} \in \{0, 1, -1\}^{k-1} : \mathbf{c} \text{ has at least } r \text{ ones}\}$, then the latter can be written as

$$\sum_{\mathbf{c} \in \mathcal{C}_{t-1}} \lambda_i \Psi_{1,\mathbf{c}} + \sum_{\mathbf{c} \in \mathcal{C}_t} \Psi_{1,\mathbf{c}} + \sum_{\mathbf{c} \in \mathcal{C}_t} (-\lambda_i - 1) \Psi_{1,\mathbf{c}} = \sum_{\mathbf{c} \in \mathcal{C}_{t-1}} \lambda_i \Psi_{1,\mathbf{c}} - \sum_{\mathbf{c} \in \mathcal{C}_t} \lambda_i \Psi_{1,\mathbf{c}}.$$

Simplifying further we obtain the following expression for $\Lambda_{i11\dots 1}$:

$$\Lambda_{i11\dots 1} = \lambda_i \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \Psi_{2,\mathbf{b}}. \quad (3.3)$$

Next we establish the upper bound on $\Lambda_{i11\dots 1}$:

$$\begin{aligned} d_1 = \Lambda_{11\dots 1} &= \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} d \Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=-1} (n-d-1) \Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=0} \Psi_{2,\mathbf{b}} \geq \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} d \Psi_{2,\mathbf{b}} \\ &\geq \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} d \Psi_{2,\mathbf{b}} = \frac{d}{\lambda_i} \Lambda_{i11\dots 1}. \end{aligned}$$

In other words, $\Lambda_{i11\dots1} \leq \frac{\lambda_i}{d} d_1$. As for a lower bound on $\Lambda_{i11\dots1}$,

$$\Lambda_{i11\dots1} = \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \lambda_i \Psi_{2,\mathbf{b}} = \frac{\lambda_i}{d} \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \Psi_{1,\mathbf{b}} \geq \frac{\lambda_i}{d} \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones, and no zeros}}} \Psi_{1,\mathbf{b}}.$$

If we assume that $d \leq n - d - 1$, then $\Psi_{1,\mathbf{b}}$ is maximized when \mathbf{b} has the maximal number of -1 's possible, or in other words, \mathbf{b} has exactly t ones and no zeros. The total fraction of such vectors out of \mathcal{B} is trivially at least $1/3^k$. Recalling that $d_1 = \sum_{\mathbf{b} \in \mathcal{B}} \Psi_{1,\mathbf{b}}$, we get

$$\Lambda_{i11\dots1} \geq \frac{1}{3^k} \cdot \frac{\lambda_i}{d} \cdot d_1.$$

Similar arguments to the ones above apply also to the case in which more than one entry i_j is greater than 1. For example it follows by the analysis above that $\Lambda_{i_1 i_2 11\dots1}$, for both $i_1, i_2 > 1$, is smaller than $\Lambda_{i11\dots1}$ by a factor of λ_j/d . This together with Eq. (3.3) imply that $\Lambda = \Lambda_{i^* 11\dots1}$ for i^* s.t. $\lambda_{i^*} = \lambda(G)$ (where i^* is either 2 or n).

To conclude, we obtained the following estimates on Λ :

$$\frac{1}{3^k} \cdot \frac{\lambda}{d} \cdot d_1 \leq \Lambda \leq \frac{\lambda}{d} \cdot d_1.$$

This completes the proof of the theorem.

4 Proof of Theorem 2.6

The proof for the bi-partite product is more complicated than the non-bipartite case. The main reason is that while there exists in the literature an exact characterization of the spectrum of $GP_{k,t}(G)$, we do not have this for $BGP_{k,t}(G)$. One of the technical contributions of this paper is obtaining such a characterization, which also enables us to prove the bound on Λ in Theorem 2.6. To understand the spectrum of $BGP_{k,t}(G)$ we embed it in a larger graph, whose spectrum we can analyze, and then infer back the spectrum of $BGP_{k,t}(G)$.

Next we define the graph in which we embed $BGP_{k,t}(G)$, and state a theorem that characterizes its spectrum. Recall the definition of the **template** of an element $\mathbf{a} = (a_1, \dots, a_k) \in (X \cup Y)^k$ given in Section 2.1: the vector $\boldsymbol{\tau} \in \{\mathcal{X}, \mathcal{Y}\}^k$ that satisfies $\boldsymbol{\tau}_i = \mathcal{X}$ iff $a_i \in X$. We are now ready to define the *shuffled bi-partite graph product* in which we embed $BGP_{k,t}(G)$.

Definition 4.1. *If $G = (X, Y, E)$ is a simple bipartite graph, define the **shuffled bi-partite graph product** $SGP_{k,t}(G)$ to be the graph with vertex set $(X \cup Y)^k$, and an edge between (a_1, \dots, a_k) and (b_1, \dots, b_k) if their templates are complements of each other and there are at least t pairs (a_i, b_i) that share an edge in G .*

In other words, the graph $SGP_{k,t}(G)$ is composed of the disjoint union of 2^{k-1} graphs $BGP_{k,t,\boldsymbol{\tau}}$ for all template pairs $\boldsymbol{\tau}, \boldsymbol{\tau}^c$, amongst which is $BGP_{k,t}(G)$ itself. For example, let $G = (X, Y, E)$ be defined with $X = \{x_1, x_2\}, Y = \{y_1, y_2\}$, with only $(x_1, y_1) \in E$ and let $k = 2, t = 1$. In this case, $SGP_{k,t}(G)$ consists of the disjoint union of the two graphs $BGP_{k,t,\boldsymbol{\tau}}$ and $BGP_{k,t,\boldsymbol{\tau}'}$ for $\boldsymbol{\tau} = \mathcal{X}^2$ and $\boldsymbol{\tau}' = \mathcal{X}\mathcal{Y}$ (a total of $16 = |X \cup Y|^2$ vertices); while $BGP_{k,t}(G) = BGP_{k,t,\boldsymbol{\tau}}$ (a total of $8 = |X|^2 + |Y|^2$ vertices).

Notice that $SGP_{k,t}(G)$ differs from $GP_{k,t}(G)$ (clearly it's a subgraph of $GP_{k,t}(G)$, yet not an induced one). In our example, the two vertices (x_1, y_1) and (y_1, y_2) are connected in $GP_{k,t}(G)$ but not connected in $SGP_{k,t}(G)$ (as the two templates corresponding to (x_1, y_1) and (y_1, y_2) are not complements of each other). Therefore we have the inclusion relation $BGP_{k,t}(G) \subseteq SGP_{k,t}(G) \subseteq GP_{k,t}(G)$.

The next theorem gives an exact characterization of the spectrum of $SGP_{k,t}(G)$. In what follows, for two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, $\mathbf{u} \otimes \mathbf{v}$ stands for the standard tensor product of two vectors.

Theorem 4.2. *Let $G = (X, Y, E)$ be a connected bi-partite n -vertex d -regular graph with eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n = -d$, and corresponding eigenvectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$, forming an orthonormal basis of \mathbb{R}^n . Define $\lambda_1^* = n/2 - d$, $\lambda_n^* = -n/2 + d$, and $\lambda_i^* = -\lambda_i$ for every $i \neq 1, n$. Let \mathcal{B} be the set of all vectors in $\{-1, 1\}^k$ having at least t 1-entries. Then $SGP_{k,t}(G)$ has a spectrum consisting of eigenvalues*

$$\Lambda_{i_1 i_2 \dots i_k} = \sum_{\mathbf{b} \in \mathcal{B}} \prod_{j=1}^k \left(\frac{1+\mathbf{b}_j}{2} \lambda_{i_j} + \frac{1-\mathbf{b}_j}{2} \lambda_{i_j}^* \right), \quad i_1, \dots, i_k \in \{1, \dots, n\}. \quad (4.1)$$

The corresponding eigenvector is $\mathbf{w}_{i_1 i_2 \dots i_k} = \mathbf{u}_{i_1} \otimes \mathbf{u}_{i_2} \otimes \dots \otimes \mathbf{u}_{i_k}$.

As we mentioned in the introduction, $GP_{k,t}(G)$ is a special case of the graph product defined in [12]. Theorem 4.2 is an analogue of Theorem 3.1 from [12]. We note that the proof of Theorem 4.2, given in Section 4.1, is similar in nature to the proof of Theorem 3.1. The following proposition gives a different view on the eigenvalues of $SGP_{k,t}(G)$.

Proposition 4.3. *Let $SGP_{k,t}(G)$ be the shuffled graph product of an n -vertex bi-partite d -regular graph $G = (X, Y, E)$. Let $\Lambda_1, \dots, \Lambda_{n^k}$ be the eigenvalues of $SGP_{k,t}(G)$. Define*

$$I = \{i : \Lambda_i = \Lambda_{i_1 i_2 \dots i_k} \text{ s.t. every } i_j \in \{1, n\}\}, \quad \Lambda = \max_{i \notin I} |\Lambda_i|. \quad (4.2)$$

Then the third largest eigenvalue of $BGP_{k,t}(G)$ is exactly Λ .

The proof of Proposition 4.3 is given at the end of this section and uses Theorem 2.8. We are now ready to prove Theorem 2.6.

Proof of Theorem 2.6: Proposition 4.3 gives us the exact characterization of Λ . This characterization is analogous to that of Λ in $GP_{k,t}(G)$: Λ is the maximal among all $\Lambda_{i_1 i_2 \dots i_k}$ in which at least one i_j does not correspond to any of the largest eigenvalues of G (in absolute value), which are $-d$ or d in the bi-partite case (i.e. $i_j \neq 1, n$). Therefore, we can use very similar arguments to the ones in the proof of Theorem 2.4 to obtain exactly the same lower and upper bounds on Λ . For the sake of completeness we give the full calculation.

Let us start by computing the value of such eigenvalues $\Lambda_{i_1 i_2 \dots i_k}$ where $i_1 \notin \{1, n\}$ and $i_s \in \{1, n\}$ for every $s = 2, \dots, k$, and then explain why indeed such eigenvalues give Λ . By Theorem 4.2, we may assume w.l.o.g. that $i_j = 1$ for all $j \geq 2$. The choice of $i_j = 1$ or $i_j = n$ only effects the sign of $\Lambda_{i_1 i_2 \dots i_k}$, which will not matter in our case since we are only interested in the absolute value. By symmetry it also doesn't matter which index is not 1 or n , so might as well assume it is the first. For the rest of the proof, we think of i_1, i_2, \dots, i_k as fixed.

For a vector \mathbf{b} , let $\#\mathbf{b}$ be its length. Denote by $\Psi_{s,\mathbf{b}} = \prod_{j=s}^{\#\mathbf{b}} \left(\frac{1+\mathbf{b}_j}{2} d + \frac{1-\mathbf{b}_j}{2} \left(\frac{n}{2} - d \right) \right)$. Using these notations and Eq. (4.1) we obtain

$$\Lambda_{i_1 i_2 \dots i_k} = \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} \lambda_i \Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=-1} -\lambda_i \Psi_{2,\mathbf{b}}.$$

Define $\mathcal{C}_r = \{\mathbf{c} \in \{1, -1\}^{k-1} : \mathbf{c} \text{ has at least } r \text{ ones}\}$, then the latter can be written as $\sum_{\mathbf{c} \in \mathcal{C}_{t-1}} \lambda_i \Psi_{1,\mathbf{c}} + \sum_{\mathbf{c} \in \mathcal{C}_t} -\lambda_i \Psi_{1,\mathbf{c}}$. Simplifying further we obtain the following expression for $\Lambda_{i11\dots1}$:

$$\Lambda_{i11\dots1} = \lambda_i \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \Psi_{2,\mathbf{b}}. \quad (4.3)$$

Next we establish the upper bound on $\Lambda_{i11\dots1}$:

$$\begin{aligned} d_2 = \Lambda_{11\dots1} &= \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} d\Psi_{2,\mathbf{b}} + \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=-1} \left(\frac{n}{2} - d\right) \Psi_{2,\mathbf{b}} \geq \sum_{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1} d\Psi_{2,\mathbf{b}} \\ &\geq \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} d\Psi_{2,\mathbf{b}} = \frac{d}{\lambda_i} \Lambda_{i11\dots1}. \end{aligned}$$

In other words, $\Lambda_{i11\dots1} \leq \frac{\lambda_i}{d} d_2$. As for a lower bound on $\Lambda_{i11\dots1}$,

$$\Lambda_{i11\dots1} = \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \lambda_i \Psi_{2,\mathbf{b}} = \frac{\lambda_i}{d} \sum_{\substack{\mathbf{b} \in \mathcal{B}, \mathbf{b}_1=1 \\ \mathbf{b} \text{ has exactly } t \text{ ones}}} \Psi_{1,\mathbf{b}}.$$

If we assume that $d \leq \frac{n}{2} - d$, then $\Psi_{1,\mathbf{b}}$ is maximized when \mathbf{b} has the maximal number of -1 's possible, or in other words, \mathbf{b} has exactly t ones. The total fraction of such vectors out of \mathcal{B} is trivially at least $1/2^k$. Recalling that $d_2 = \sum_{\mathbf{b} \in \mathcal{B}} \Psi_{1,\mathbf{b}}$, we get

$$\Lambda_{i11\dots1} \geq \frac{\lambda_i}{d} \cdot d_2 \cdot \frac{1}{2^k}.$$

Similar arguments to the ones above imply for example that $\Lambda_{ij11\dots1}$, for both $i, j \notin \{1, n\}$, is smaller than $\Lambda_{i11\dots1}$. Thus, as before, together with Eq. (4.3) we obtain that $\Lambda = \Lambda_{i^*11\dots1}$ for i^* s.t. $\lambda_{i^*} = \lambda(G)$.

To conclude, we obtained the following estimates on Λ :

$$\frac{1}{2^k} \cdot \frac{\lambda}{d} \cdot d_2 \leq \Lambda \leq \frac{\lambda}{d} \cdot d_2.$$

This completes the proof of the theorem.

4.1 Proof of Theorem 4.2

Our proof strategy, which follows the outline of the proof of Theorem 3 in [12], is as follows: We are going to identify the adjacency matrix \hat{A} of the graph $SGP_{k,t}(G)$, and then analyze its spectrum, and show that its eigenvectors form an orthonormal basis for \mathbb{R}^{n^k} .

We start with a few notations and definitions. Let A be the adjacency matrix of the d -regular bi-partite graph $G = (X, Y, E)$, and let $d = \lambda_1, \dots, \lambda_n = -d$ be its eigenvalues with corresponding eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n$, which we may assume form an orthonormal basis for \mathbb{R}^n . Let C be the following block matrix: $C_{ij} = 0$ if $i, j \in X$ or if $i, j \in Y$, and $C_{ij} = 1$ otherwise (in words, C is the adjacency matrix of the complete bi-partite graph). Define the bi-partite complement of G , denoted

by \bar{G} , to be the bi-partite graph whose adjacency matrix is given by $\bar{A} = C - A$. For fixed $i \in [k]$ and $\mathbf{b} \in \mathcal{B}$, define the auxiliary matrix:

$$M_{i\mathbf{b}} = \frac{1 + \mathbf{b}_i}{2} A + \frac{1 - \mathbf{b}_i}{2} \bar{A}.$$

In words, $M_{i\mathbf{b}}$ is either the adjacency matrix of G , if $\mathbf{b}_i = 1$, or of \bar{G} if $\mathbf{b}_i = -1$.

Using these notations, we claim that the adjacency matrix of $SGP_{k,t}(G)$ is

$$\hat{A} = \sum_{\mathbf{b} \in \mathcal{B}} M_{1\mathbf{b}} \otimes M_{2\mathbf{b}} \otimes \cdots \otimes M_{k\mathbf{b}}. \quad (4.4)$$

By the definition of matrix tensor product, every entry in \hat{A} has the form $\sum_{\mathbf{b} \in \mathcal{B}} \prod_{i=1}^k (M_{i\mathbf{b}})_{h_i g_i}$, for some $\mathbf{h} = (h_1, \dots, h_k)$ and $\mathbf{g} = (g_1, \dots, g_k)$ in $(X \cup Y)^k$. We address that entry in \hat{A} by $\hat{A}_{\mathbf{h}, \mathbf{g}}$, namely

$$\hat{A}_{\mathbf{h}, \mathbf{g}} = \sum_{\mathbf{b} \in \mathcal{B}} \prod_{i=1}^k (M_{i\mathbf{b}})_{h_i g_i}. \quad (4.5)$$

From the definition of the tensor product, it follows that the rows and column of \hat{A} are indexed according to the following lexicographic order on $X \cup Y$: $x_i < y_j$ for every i, j , $x_i < x_j$ if $i < j$ and similarly $y_i < y_j$ if $i < j$. For example, the first row of \hat{A} is indexed by the vertex (x_1, \dots, x_1, x_1) , the second row by (x_1, \dots, x_1, x_2) , and so on.

Now we can show that $\hat{A}_{\mathbf{h}, \mathbf{g}} = 1$ if \mathbf{h} and \mathbf{g} share an edge in $SGP_{k,t}(G)$, and $\hat{A}_{\mathbf{h}, \mathbf{g}} = 0$ otherwise. Let us start with the easier case, where the templates of \mathbf{h} and \mathbf{g} are not complements of each other. We need to show that $\hat{A}_{\mathbf{h}, \mathbf{g}} = 0$ as by the definition of $SGP_{k,t}(G)$ they cannot share an edge. Since the templates are not complementary, there exists an index i s.t. both h_i and g_i belong to, say, X . Therefore, no matter what $\mathbf{b} \in \mathcal{B}$ we choose, $(M_{i\mathbf{b}})_{h_i g_i} = 0$, since (h_i, g_i) is not an edge of G nor of \bar{G} . Therefore the product in Eq. (4.5) is always zero. Next assume that \mathbf{h} and \mathbf{g} have complementary templates. Let \mathbf{b} be the vector whose entries are $\mathbf{b}_i = 1$ if $(h_i, g_i) \in E(G)$ and $\mathbf{b}_i = -1$ otherwise. If indeed \mathbf{h} and \mathbf{g} share an edge in $SGP_{k,t}(G)$, then for this specific \mathbf{b} , the value $(M_{i\mathbf{b}})_{h_i g_i}$ is 1 for every $i \in [k]$, since by definition the vector \mathbf{b} “selects” the correct matrix A or \bar{A} . For all other \mathbf{b} ’s, some $(M_{i\mathbf{b}})_{h_i g_i}$ is going to be zero, as for some index i the “wrong” adjacency matrix is going to be chosen. All in all, if \mathbf{h} and \mathbf{g} share an edge in $SGP_{k,t}(G)$, then $\hat{A}_{\mathbf{h}, \mathbf{g}} = 1$. Conversely, if \mathbf{h} and \mathbf{g} don’t share an edge in $SGP_{k,t}(G)$, then for every $\mathbf{b} \in \mathcal{B}$, there must be some index i s.t. $\mathbf{b}_i = 1$ but $(h_i, g_i) \notin E$ (otherwise, there are at least t pairs (h_i, g_i) that share an edge in G , and by the definition of $SGP_{k,t}(G)$ we should have placed an edge between \mathbf{h} and \mathbf{g}). In this case $A_{h_i g_i} = 0$ and also by definition $M_{i\mathbf{b}} = A$. Thus, the product in Eq. (4.5) is zeroed.

Next let us show that $\mathbf{w}_i = \mathbf{w}_{i_1 i_2 \dots i_k} = \mathbf{u}_{i_1} \otimes \mathbf{u}_{i_2} \cdots \otimes \mathbf{u}_{i_k}$ is an eigenvector of \hat{A} , with eigenvalue $\Lambda_i = \Lambda_{i_1 \dots i_k}$, as defined in Eq. (4.1). To that end, we first show that the eigenvalues of \bar{A} are exactly $\lambda_1^*, \lambda_2^*, \dots, \lambda_n^*$ as defined in Theorem 4.2, with eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n$. This follows immediately from the definition of $\bar{A} = C - A$, and the simple observation that $C\mathbf{u}_i = 0$ for every $i \neq 1, n$ (since $\mathbf{u}_i \perp \mathbf{u}_1, \mathbf{u}_n$, it is also perpendicular to every row of C which is simply $0.5(\mathbf{u}_1 + \mathbf{u}_n)$ or $0.5(\mathbf{u}_1 - \mathbf{u}_n)$). To see why Λ_i is an eigenvector of \hat{A} , use the following fact about tensor products: for two matrices P, Q and two vectors \mathbf{u}, \mathbf{v} : $(P \otimes Q)(\mathbf{u} \otimes \mathbf{v}) = (P\mathbf{u}) \otimes (Q\mathbf{v})$.

4.2 Proof of Proposition 4.3

Let $G = (X, Y, E)$ be a connected bi-partite d -regular n -vertex graph, and let $SGP_{k,t}(G)$ be its shuffled bi-partite graph product with vertex set $V = (X \cup Y)^k$. Set $N = 2(n/2)^k$ and $M = n^k$. Recall that the different $BGP_{k,t,\tau}(G)$'s are cospectral (Theorem 2.8), and let us denote by $\Psi_1, \Psi_2, \dots, \Psi_N$ their set of eigenvalues with corresponding eigenvectors $\mathbf{v}_1^\tau, \dots, \mathbf{v}_N^\tau$. Similarly, let $\Lambda_1, \Lambda_2, \dots, \Lambda_M$, be the eigenvalues of $SGP_{k,t}(G)$.

The first observation that we make is that the spectrum of $SGP_{k,t}(G)$ consists of the eigenvalues Ψ_1, \dots, Ψ_N , each with multiplicity $2^k/2$. This clearly accounts for the entire spectrum of $SGP_{k,t}(G)$ since $N \cdot 2^k/2 = M$. To see this, for every template pair τ, τ^c , embed the vector \mathbf{v}_i^τ in \mathbb{R}^M by padding with zeros for all entries that do not correspond to τ, τ^c . The padded vector is an eigenvector of $SGP_{k,t}(G)$ with eigenvalue Ψ_i . In this way we obtain $2^k/2$ linearly independent eigenvectors of $SGP_{k,t}(G)$ for every eigenvalue Ψ_i .

The next step is to figure out the eigenvalues Λ_i that correspond to the set I defined in Eq.(4.2). We shall use Theorem 4.2 to carry out this task. Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of G with corresponding eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n$. Since G is bi-partite we have $\mathbf{u}_n = (\mathbf{1}_{n/2}, -\mathbf{1}_{n/2})$, as well as $\mathbf{u}_1 = \mathbf{1}_n$. Recall the definition of Λ_i given in Eq. (4.1),

$$\Lambda_i = \Lambda_{i_1 \dots i_k} = \sum_{\mathbf{b} \in \mathcal{B}} \prod_{j=1}^k \left(\frac{1 + \mathbf{b}_j}{2} \lambda_{i_j} + \frac{1 - \mathbf{b}_j}{2} \lambda_{i_j}^* \right). \quad (4.6)$$

By the definition of the set I , every i_j is either 1 or n . Fix $\mathbf{b} \in \mathcal{B}$ with exactly t ones. Let p denote the total number of indices j s.t. $i_j = n$ in Λ_i , and let r stand for the number of j 's s.t. $i_j = n$ and $\mathbf{b}_j = 1$. Substituting λ_{i_j} and $\lambda_{i_j}^*$ according to their definition in Theorem 4.2, the product in Eq. (4.6) equals exactly

$$d^{t-r}(-d)^r \left(-\frac{n}{2} + d \right)^{p-r} \left(\frac{n}{2} - d \right)^{k-t-(p-r)} = (-1)^p d^t \left(\frac{n}{2} - d \right)^{k-t}.$$

Summing over all vectors $\mathbf{b} \in \mathcal{B}$ we get

$$\Lambda_i = (-1)^p \sum_{t'=t}^k d^{t'} \left(\frac{n}{2} - d \right)^{k-t'} \binom{k}{t'}. \quad (4.7)$$

Recalling the value of the degree d_2 of $BGP_{k,t}(G)$ in Eq. (2.1), we get that Eq. (4.7) is exactly $(-1)^p d_2$. The set I is of size 2^k , and exactly half of the choices $i_1 \dots i_k$ have an odd number of j 's s.t. $i_j = n$. Therefore the set $\{\Lambda_i : i \in I\}$ contains $d_2 2^k/2$ times, and $-d_2 2^k/2$ times. Since $BGP_{k,t,\tau}(G)$ is bi-partite, we have $\Psi_1 = d_2$ and $\Psi_N = -d_2$. The latter, together with the previous observation imply that the set $\{\Lambda_i : i \notin I\} = \{\Psi_2, \Psi_3, \dots, \Psi_{N-1}\}$. In turn, $\Lambda = \max_{i \notin I} |\Lambda_i| = \max\{\Psi_2, \Psi_{N-1}\}$, which is the third largest eigenvalue of $BGP_{k,t}(G)$ in absolute value. This completes the proof of Proposition 4.3.

5 Proof of Theorem 2.8

Before we approach the actual proof we need a few preliminary results and definitions. The following well known lemma gives a necessary and sufficient condition for two graphs to be cospectral (see for example [8], Lemma 2.1).

Lemma 5.1. Let G_1 and G_2 be two n -vertex graphs, with adjacency matrices A_1 and A_2 respectively. The two graphs G_1 and G_2 are cospectral iff $\text{tr}(A_1^\ell) = \text{tr}(A_2^\ell)$ for every $\ell \geq 1$.

Let us state the lemma in words. Recall that the i^{th} diagonal entry of the ℓ^{th} power of the adjacency matrix of a graph is the number of closed paths that contain vertex i in the graph. Hence the trace of the ℓ^{th} power is the sum of closed paths of length ℓ with over counting. In words, Lemma 5.1 states that two graphs are co-spectral, if these quantities are the same in both for every ℓ . We shall prove a somewhat stronger claim, that in particular implies the latter. For an ordered set $p = (a_1, a_2, \dots, a_\ell, a_{\ell+1})$ of vertices in $(X \cup Y)^k$, define its connectivity vector $\mathbf{c}(p)$ according to the following rule: for every index i s.t. $a_i = (u_1, \dots, u_k)$ and $a_{i+1} = (v_1, \dots, v_k)$ that share exactly q edges $(u_s, v_s) \in E(G)$, we set $\mathbf{c}_i = q$. In particular, p is a cycle in $BGP_{k,t,\tau}$ iff $\mathbf{c}_i(p) \geq t$ for every i . We say that an ordered set $(a_1, a_2, \dots, a_\ell, a_{\ell+1})$ is *circular and alternating* if $a_{\ell+1} = a_1$, and every two adjacent vertices in the set belong to complementary templates.

We shall prove the following proposition regarding the set of connectivity vectors that belong to different templates. As a simple corollary we get the cospectrality result. For simplicity, we treat all sets of vertices as ordered sets. This incurs over counting, which will not matter, as it cancels out. Also observe that two different sets of vertices may have the same connectivity vector. We treat the two vectors as different vectors. In other words, the sets \mathcal{C}_τ defined below should be treated as *multi-sets*.

Proposition 5.2. Fix ℓ , and let \mathcal{C}_τ be the multi-set of connectivity vectors of all circular and alternating ℓ -vertex sets in $BGP_{k,t,\tau}$, and similarly define $\mathcal{C}_{\tau'}$. Then $\mathcal{C}_{\tau'} = \mathcal{C}_\tau$ for any pair of templates τ, τ' and any $\ell \geq 0$.

Before proving the proposition, we establish the following auxiliary result concerning bi-partite d -regular graphs $G = (X \cup Y, E)$. We say that an ordered set $(v_1, v_2, \dots, v_\ell, v_1)$ obeys a vector $\mathbf{u} \in \{0, 1\}^\ell$ if (v_i, v_{i+1}) is an edge of G iff $\mathbf{u}_i = 1$ (for $i = \ell$ we consider the edge between v_ℓ and v_1). Let $\Psi_{\mathbf{u},X}$ be the set of all alternating closed paths in G that obey \mathbf{u} and $v_1 \in X$. Similarly define $\Psi_{\mathbf{u},Y}$.

Lemma 5.3. Let $G = (X \cup Y, E)$ be a d -regular n -vertex bi-partite graph. For every $\ell \geq 0$ and every vector $\mathbf{u} \in \{0, 1\}^\ell$, we have $|\Psi_{\mathbf{u},Y}| = |\Psi_{\mathbf{u},X}|$.

Proof. If ℓ is odd, then clearly $|\Psi_{\mathbf{u},Y}| = |\Psi_{\mathbf{u},X}| = 0$, since G is bi-partite. Now consider an even ℓ . Let C be the adjacency matrix of the complete bi-partite n -vertex graph, and A be the adjacency matrix of G . Recall our definition of the bi-partite complement of G , denoted by \bar{G} , which is the bi-partite graph whose adjacency matrix is given by $\bar{A} = C - A$. Let I_1 be the $n \times n$ matrix, whose first $n/2$ diagonal entries are 1, and all other entries are 0. Similarly define I_2 w.r.t. to the last $n/2$ diagonal entries. Using these definitions we claim that

$$|\Psi_{\mathbf{u},X}| = \text{tr} \left(I_1 \cdot \prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A}) \right), \quad |\Psi_{\mathbf{u},Y}| = \text{tr} \left(I_2 \cdot \prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A}) \right). \quad (5.1)$$

The j^{th} diagonal entry in the matrix product $\prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A})$ is the total number of closed paths that start at x_j (if $j \leq n/2$, or $y_{j-n/2}$ if $j > n/2$) and obey \mathbf{u} . If $\mathbf{u} = 1^\ell$ for example, then we get the standard interpretation of taking powers of the adjacency matrix. Multiplying by I_1 (or I_2) simply extracts the relevant paths: starting at X or Y (the trace of a matrix product obeys the rule $\text{tr}(A^T B) = \sum_{i,j} A_{ij} B_{ij}$).

The product $\prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A})$ defines a polynomial in A (replace \bar{A} by $C - A$). When taking the actual products, we get two types of terms: we get A^ℓ , or terms of the form $(-1)^{s_2+s_4+\dots} A^{s_1} C^{s_2} A^{s_3} \dots$ for some natural numbers s_i . As A and C commute, one can readily verify that the second type of terms simplify to $b \cdot (J - C)$, where J is the all-one matrix, and $b = d^{s_1+s_3+\dots} \left(\frac{n}{2}\right)^{-1+s_2+s_4+\dots}$. Let us denote this polynomial by $P = A^\ell + T$, where T is the sum of all type-two terms. Using the linearity of the trace operator, Eq. (5.1) can be rewritten as $\text{tr} \left(I_1 \cdot \prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A}) \right) = \text{tr}(I_1 A^\ell) + \text{tr}(I_1 T)$. By our characterization of T , it holds that $\text{tr}(I_1 T) = \text{tr}(I_2 T)$. As for $\text{tr}(I_1 A^\ell)$, this is the sum over all i of the number of closed paths of length ℓ that start at x_i . The same is true for $\text{tr}(I_2 A^\ell)$ just with y_i . However there is a natural bijection between the two sets: map $x_{i_1} - y_{i_1} - x_{i_2} - \dots$ to $y_{i_1} - x_{i_2} - \dots$. To conclude, we got

$$\begin{aligned} |\Psi_{\mathbf{u},X}| &= \text{tr} \left(I_1 \cdot \prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A}) \right) = \\ &= \text{tr}(I_1 T) + \text{tr}(I_1 A^\ell) = \text{tr}(I_2 T) + \text{tr}(I_2 A^\ell) = \text{tr} \left(I_2 \cdot \prod_{i=1}^{\ell} (\mathbf{u}_i A + (1 - \mathbf{u}_i) \bar{A}) \right) = |\Psi_{\mathbf{u},Y}|. \end{aligned}$$

■

We are now ready to prove Proposition 5.2

Proof.(Proposition 5.2) We fix the parameter ℓ and prove via induction on k . The base case, $k = 1$, is true in a vacuous way since for $k = 1$ there is only one template pair. Assume the claim is true up to $k - 1$ and let us prove it for k . Fix a template τ of length k , and let $\hat{\tau}$ be its $(k - 1)$ -prefix. The graph $BGP_{k,t,\tau}$ is obtained from $BGP_{k-1,t,\hat{\tau}}$ by extending each $(k - 1)$ -tuple with a vertex from $X \cup Y$ (according to τ) and updating the edge set accordingly (existing edges remain, but new edges may appear).

Let $\mathbf{c} \in \mathcal{C}_{\hat{\tau}}^{(k-1)}$ be a connectivity vector corresponding to the (circular and alternating) ordered set $\mathbf{a} = (a_1, a_2 \dots, a_\ell, a_{\ell+1})$ with each a_i in $BGP_{k-1,t,\hat{\tau}}$. Let $\mathbf{u} \in \{0, 1\}^\ell$. An update of \mathbf{a} by \mathbf{u} is an *extended* version $\mathbf{a}^{(k)}$ of \mathbf{a} in which one appends to each $a_i \in (X \cup Y)^{k-1}$ a vertex v_i in $X \cup Y$ (according to the template τ) such that $v = (v_1, v_2, \dots, v_\ell, v_{\ell+1})$ in an alternating closed path that obeys \mathbf{u} , i.e., v is either in $\Psi_{\mathbf{u},X}$ or $\Psi_{\mathbf{u},Y}$. For any such extension of $\mathbf{a}^{(k)} = (a_1^{(k)}, \dots, a_\ell^{(k)}, a_{\ell+1}^{(k)})$ (which is circular and alternating and consists of k -tuples $a_i^{(k)} \in (X \cup Y)^k$ of $BGP_{k,t,\tau}$), the corresponding connectivity vector is exactly $\mathbf{c} + \mathbf{u}$. The set \mathcal{C}_τ consists of the multi-set of connectivity vectors obtained by this process starting from any \mathbf{a} and any \mathbf{u} as above.

To prove the induction step, it suffices to show that for any connectivity vector $\mathbf{c} \in \mathcal{C}_{\hat{\tau}}^{(k-1)}$ corresponding to a circular and alternating $\mathbf{a} = (a_1, a_2 \dots, a_\ell, a_{\ell+1})$ in $BGP_{k-1,t,\hat{\tau}}$ and any vector $\mathbf{u} \in \{0, 1\}^\ell$, the number extensions $\mathbf{a}^{(k)}$ of \mathbf{a} by \mathbf{u} does not depend on \mathbf{a} , τ or $\hat{\tau}$.

The latter follows from the fact that the number of alternating closed paths $v = (v_1, v_2, \dots, v_\ell, v_{\ell+1})$ that obey \mathbf{u} is determined completely by $G = (X, Y, E)$ and is independent of \mathbf{a} and $\hat{\tau}$. However, the value at hand may potentially depend on τ only through the question whether $v_1 \in X$ or $v_1 \in Y$. For $v_1 \in X$ the number of alternating closed paths v is the size of $\Psi_{\mathbf{u},X}$ and for $v_1 \in Y$, the size of $\Psi_{\mathbf{u},Y}$. By Lemma 5.3 these two are equal and we may conclude our proof. ■

6 Proof of Proposition 2.3

Fix a vertex $\mathbf{x} = (x_1, \dots, x_k)$, and let us compute the number of neighbors (y_1, \dots, y_k) that \mathbf{x} has in $GP_{k,t}(G)$. Fix some $t' \geq t$, there are $\binom{k}{t'}$ ways to choose the indices j for which $(x_j, y_j) \in E(G)$. For every such j , there are d possible y_j 's. For the other x_j 's, there are $n - d$ non-neighbors to choose from (this includes x_j itself, as the graph G is simple). Now sum over all possible $t' \geq t$ to obtain

$$d_1 = \sum_{t'=t}^k \binom{k}{t'} d^{t'} (n-d)^{k-t'}.$$

Since \mathbf{x} was arbitrary, this equality holds for every \mathbf{x} , hence $GP_{k,t}(G)$ is d_1 -regular. Similar arguments give the bi-partite case, and the shuffled product, which is composed of 2^{k-1} disjoint bi-partite graphs, similar to $BGP_{k,t}(G)$.

References

- [1] N. Alon and F. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, **72** (1988), 15–19.
- [2] N. Alon and E. Lubetzky. Graph Powers, Delsarte, Hoffman, Ramsey, and Shannon. *SIAM J. Discret. Math.* **21** (2007), 329–348.
- [3] Y. Bilu, and N. Linial. Lifts, Discrepancy and Nearly Optimal Spectral Gap. *Combinatorica* **26** (2006), 495–519.
- [4] D. Cvetković, Graphs and their spectra (Grafovi i njihovi spektri) (Thesis), Univ. Beograd Publ. Elektrotehn. Fak., Ser. Mat. Fiz., No. 354No. 356, (1971) 1-50.
- [5] D. Cvetković, M. Doob, H. Sachs, Spectra of Graphs-Theory and Application, third edition, Johann Ambrosius Barth Verlag, 1995
- [6] M. Garey and D. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. *W. H. Freeman & Co.* New York, NY, USA, 1979.
- [7] S. Hedetniemi. Homomorphisms and graph automata, University of Michigan Technical Report 03105-44-T, 1966.
- [8] C. D. Godsil and B. D. McKay. Products of graphs and their spectra, University of Michigan *Combinatorial Mathematics IV, Lecture Notes in Mathematics* **560** (1976), 61–72.
- [9] C. D. Godsil and B. D. McKay. Constructing cospectral graphs. *Aequationes Math.* **25** (1982), 257–268.
- [10] M. Krivelevich and B. Sudakov. Pseudo-Random Graphs. *More sets graphs and number, Bolyai Society Mathematical Studies* **15** (2006) , 199–262.
- [11] A. Laub. Matrix Analysis for Scientists and Engineers. *Society for Industrial and Applied Mathematics Philadelphia, PA, USA*, 2004.

- [12] R. Šokarovski. A generalized direct product of graphs. *Publications De L'Institut Mathematique, nouvelle serie* **36** (1977), 267–269.
- [13] A. Thomason. Graph products and monochromatic multiplicities. *Combinatorica* **17** (1997), 125–134.
- [14] A. Thomason. Pseudorandom graphs. *Random Graphs '85, North-Holland Math. Stud.* **144** (1987), 307-331.
- [15] J. H. van Lint and J. J. Seidel . Equilateral point sets in elliptic geometry. *Indag. Math.* **28** (1966), 335-348.
- [16] P. Weichsel. The Kronecker product of graphs. *Proceedings of the American Mathematical Society* **13** (1962), 47-52.